arXiv:1703.04529v1 [cs.LG] 13 Mar 2017

Task-based End-to-end Model Learning

Priya L. Donti [1 2] **Brandon Amos** [1] **J. Zico Kolter** [1]

Abstract

As machine learning techniques have become more ubiquitous, it has become common to see machine learning prediction algorithms operating within some larger process. However, the criteria by which we train machine learning algorithms often differ from the ultimate criteria on which we evaluate them. This paper proposes an end-to-end approach for learning probabilistic machine learning models within the context of stochastic programming, in a manner that directly captures the ultimate task-based objective for which they will be used. We then present two experimental evaluations of the proposed approach, one as applied to a generic inventory stock problem and the second to a real-world electrical grid scheduling task. In both cases, we show that the proposed approach can outperform both a traditional modeling approach *and* a purely black-box policy optimization approach.

1. Introduction

As machine learning techniques have become more ubiquitous, it has become common to see machine learning prediction algorithms operating within some larger process. For instance, instead of merely classifying images in a standalone setting, one may want to use these classifications within planning and control tasks such as those that arise in autonomous driving. Similarly, when we use a probabilistic prediction algorithm to generate forecasts of upcoming electricity demand, we will then want to use these forecasts within the loop of a scheduling procedure that allocates generation for a power grid. It is well-understood in practice that in these settings, the criteria by which we train machine learning algorithms differ from the ultimate criteria on which we evaluate them. Thus, while machine learning algorithms may be trained to maximize the accuracy or log likelihood of the data, the corresponding loss functions may be only indirectly correlated with the task we ultimately care about: the performance of the full "closed-loop" system on the ultimate task at hand. For example, in a driving task, mistakenly classifying a pedestrian as a tree would have very different consequences than misclassifying a garbage can as a tree. As this example suggests, we instead may want to approximate the ultimate task-based "true loss" by carefully weighting or calibrating the training set on different classification costs.

This paper considers an end-to-end approach for learning probabilistic machine learning models, in a manner that directly captures the ultimate task-based objective for which they will be used. Formally, we consider probabilistic models in the context of stochastic programming, a setting where the goal is to minimize some expected cost over the models' probabilistic predictions, subject to some (potentially also probabilistic) constraints. As mentioned above, it is common to approach these problems in a two-stage fashion: to first fit a predictive model to observed data by minimizing some criterion such as the negative log likelihood of the model, and then to use these models to compute or approximate the necessary expectations in the stochastic programming setting. While this procedure can work well in many instances, it ignores the fact that the true cost of the system (the optimization objective evaluated on *actual* instantiations in the real world) may benefit from a model that actually attains worse overall likelihood, but which makes more accurate predictions over certain manifolds of the underlying space.

As such, we propose to train a probabilistic model not (solely) for predictive accuracy, but so that–when it is later used within the loop of a stochastic programming procedure–it produces solutions that minimize the ultimate task-based loss. This formulation may seem somewhat counterintuitive, given that a "perfect" predictive model would of course also be the optimal model to use within a stochastic programming framework. However, the reality that all models *do* make errors illustrates that we should indeed look to a final task-based objective to determine the proper tradeoffs between these errors within a machine learning setting. This paper effectively proposes one way to evaluate task-based tradeoffs in a fully automated fash-

[1]Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA [2]Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA. Correspondence to: Priya L. Donti <pdonti@cs.cmu.edu>, J. Zico Kolter <zkolter@cs.cmu.edu>.

Under review by the International Conference on Machine Learning (ICML) 2017. *Copyright 2017 by the author(s).*

ion, by computing derivatives through the solution to the stochastic programming problem in a manner that can improve the underlying model.

We begin by presenting background material and related work in areas spanning stochastic programming, end-to-end training, optimizing alternative loss functions, and the classic generative/discriminative tradeoff in machine learning. We then present our approach within the formal context of stochastic programming, and give a generic method for propagating task loss through these problems in a manner that can update the models. We present two experimental evaluations of the proposed approach, one as applied to a generic inventory stock problem and the second to a real-world electrical grid scheduling task. In both cases we show that the proposed approach can outperform both a traditional modeling approach *and* a purely black-box policy optimization approach.

2. Background and Related Work

Stochastic programming Stochastic programming is a method for making decisions under uncertainty by modeling or optimizing objectives governed by a random process. It has applications in many domains such as energy (Wallace & Fleten, 2003), finance (Ziemba & Vickson, 2006), and manufacturing (Buzacott & Shanthikumar, 1993), where the underlying probability distributions are either known or can be estimated. As presented by Shapiro & Philpott (2007), common considerations in stochastic programming include how to best model or approximate the underlying random variable, how to then solve the resulting optimization problem, and how to assess the quality of the resulting (approximate) solution.

In cases where the underlying probability distribution is known but the objective cannot be solved analytically, it is common to use Monte Carlo sample average approximation methods, which draw multiple iid samples from the underlying probability distribution and then use deterministic optimization methods (Linderoth et al., 2006). In cases where the underlying distribution is not known, it is common to learn or estimate some model from observed samples, either directly using samples from the real environment in these programs (Rockafellar & Wets, 1991) or via learning a probabilistic model.

End-to-end training Recent years have seen a dramatic increase in the number of systems building on so-called "end-to-end" learning. Generally speaking, this term refers to systems where the end goal of the machine learning process is directly predicted from the raw inputs (e.g. LeCun et al., 2005; Thomas et al., 2006). In the context of deep learning systems, the term now traditionally refers to architectures where there is no, for example, explicit encoding of hand-tuned features on the data, but the system directly predicts what the image, text, etc. is from the raw inputs (Wang et al., 2011; He et al., 2016; Wang et al., 2012; Graves & Jaitly, 2014; Amodei et al., 2015). The context in which we use the term end-to-end is similar, but slightly more in line with its older usage: instead of (just) attempting to learn an output (with known and typically straightforward loss functions), we are specifically attempting to learn based upon an end-to-end *task* that the user is ultimately trying to accomplish. Conceptually, we believe that this is a beneficial concept to add to the notion of end-to-end learning, not just for describing an end-to-end process in terms of mapping inputs from outputs, but for describing the entire closed-loop performance of the system as evaluated on the real task at hand.

Also highly related to our work are recent efforts in end-to-end policy learning (Levine et al., 2016), and recent work in using value iteration effectively as an optimization procedure in similar networks (Tamar et al., 2016). This line of work fits more with the "pure" end-to-end approach we discuss below, where models are eschewed for pure function approximation methods, but conceptually the approaches have similar motivations in modifying typical policies to address some task directly. Of course, the actual methodological approaches are quite different, given the specific focus on stochastic programming as the black box of interest in our setting.

Optimizing alternative loss functions There has been a great deal of work in recent years on using machine learning procedures (including deep networks), to optimize a different loss criterion than the one "naturally" optimized by the algorithm. For example, both Stoyanov et al. (2011) and Hazan et al. (2010) proposed methods for optimizing loss criteria in structured prediction that are *different* from the inference procedure of the prediction algorithm; the work has also been recently extended to the setting of deep networks (Song et al., 2016).

The work that we have found in the literature that most closely resembles the approach we employ here is the work of Bengio (1997), which used a neural network model for predicting financial prices, but then optimized the model based upon financial returns obtained via a hedging strategy that employed the model. We view this approach as a philosophical predecessor to our own work: the core idea of both using a model but then tuning that model to adapt to a (differentiable) procedure is akin to our own methods, though the actual methodological approach is naturally quite different. Whereas Bengio (1997)'s work used a hand-crafted (but differentiable) algorithm to approximately attain some objective given a predictive model, our approach is tightly coupled to stochastic programming, where the explicit objective is to *attempt* to optimize the

desired task cost via an exact optimization routine, just in an approximate model. The notions of stochasticity, etc, are naturally thus quite different in our own work, but we do hope that our work can bring this idea back from this previous work (despite the paper being nearly 20 years old, virtually all follow-on work has focused on the financial application aspect, and not on what we feel is the core idea of using a surrogate model within a task-driven optimization procedure.

Generative and discriminative tradeoff One last point of discussion relates to the classic tradeoff between generative and discriminative models (Ng & Jordan, 2002). This relation is more at a conceptual level: just as discriminative models forgo any attempt to model the joint distribution over inputs and outputs $p(x, y)$, and focus instead on just the conditional distribution $p(y|x)$ to directly predict y, our method in some sense goes a step further. Instead of even trying to accurately represent the distribution $p(y|x)$, we solely attempt to use whatever model parameters are necessary to optimize some subsequent task loss. However, we still go through the same intermediary form of the probabilistic model, much as discriminative and generative models sometimes have pairings (e.g. naive Bayes and logistic regression), where both models use the same underlying representation but one approach merely trains based upon conditional rather than joint likelihood. We take this a step further and use the same representation but train task loss instead of either of these likelihoods.

3. End-to-end model learning in stochastic programming

We begin by formally defining the stochastic modeling and optimization problems with which we are concerned. Let $(x \in \mathcal{X}, y \in \mathcal{Y}) \sim \mathcal{D}$ denote standard input, output pairs drawn from some (real, unknown) distribution $\mathcal{D}$. As in the standard statistical setting, we will model the conditional distribution $y|x$ using some parameterized model $p(y|x;\theta)$ for some set of parameters θ. On top of this standard setup, however, we also consider a stochastic programming problem with optimization variable $z \in \mathcal{Z}$, where we choose a policy z to minimize the expected value of our objective. Explicitly, given some inputs x, we wish to solve some generic stochastic programming problem

$$\begin{aligned} \underset{z}{\text{minimize}} \quad & \mathbf{E}[f(x, y, z)] \\ \text{subject to} \quad & \mathbf{E}[g_i(x, y, z)] \leq 0, \quad i = 1, \ldots, n_{ineq} \\ & h_i(z) = 0, \quad i = 1, \ldots, n_{eq} \end{aligned} \tag{1}$$

where the expectations in all cases are taken with respect to $y \sim p(y|x;\theta)$.[2] We denote the optimal solution to this problem as $z^\star(x;\theta)$, to emphasize that this optimal solution depends both on the observed input x and on the choice of parameters θ that determine our probabilistic model.

Although we use the optimization problem above to determine our optimal actions $z^\star$, it is actually not this objective that we ultimately care about. Rather, we want to minimize this objective as it occurs under the *true* joint distribution $x, y \sim \mathcal{D}$, where $z^\star$ is chosen by solving the problem (1) above. That is, our true objective (to which we refer as the *task loss*), is given by

$$\begin{aligned} L(\theta) = \mathbf{E}_{x,y\sim\mathcal{D}}[f(x, y, z^\star(x;\theta))] \\ + \sum_{i=1}^{n_{ineq}} I\{\mathbf{E}_{x,y\sim\mathcal{D}}[g_i(x, y, z^\star(x;\theta))] \leq 0\} \\ + \sum_{i=1}^{n_{eq}} I\{h_i(z^\star(x;\theta)) = 0\} \end{aligned} \tag{2}$$

where $I(\cdot)$ denotes the indicator function that is zero when its constraints are satisfied and infinite otherwise. The ultimate goal of this work will be to develop methods by which we can directly optimize (or at least approximate) $L(\theta)$ over θ, that is, methods that can (approximately) solve the optimization problem

$$\underset{\theta}{\text{minimize}} \quad L(\theta). \tag{3}$$

Solving this problem will require that we differentiate through the "argmin" operator $z^\star(x;\theta)$ of the stochastic programming problem; while this is not possible for all classes of optimization problems (the argmin operator may be discontinuous), in many practical cases (including e.g., cases where the function and constraints are strongly convex), we can indeed efficiently compute these gradients even in the context of constrained optimization.

3.1. Discussion

Before presenting our methodology for solving the general optimization problem in (3), we highlight the approach in contrast to two alternative existing methods: traditional model learning and model-free black-box policy optimization. In traditional ML approaches, it is common to use θ to minimize the (conditional) log likelihood of observed data under the model $p(y|x;\theta)$. This corresponds to approximately solving the optimization problem

$$\underset{\theta}{\text{minimize}} \quad \mathbf{E}_{x,y\sim\mathcal{D}}\left[-\log p(y|x;\theta)\right]. \tag{4}$$

If we then need to use the conditional distribution $y|x$ within some later optimization setting, we commonly just use the predictive model obtained from (4) directly.

[2] It is standard to presume in stochastic programming problems that equality constraints depend only on the decision variables and not the random variables, as non-trivial random equality constraints typically are not possible to satisfy.

This approach has obvious advantages, in that the model-learning phase is well-justified independent of any future use in a task. However, it is also prone to poor performance in the (completely typical) setting where the true distribution $y|x$ cannot be represented within the class of distribution parameterized by θ. In this setting, the log likelihood objective *implicitly* provides a tradeoff between model error over different regions of the input/output space, but does so in a manner that is largely opaque to the modeler, and which may ultimately *not* be the correct tradeoff to employ for any given task (e.g. a particular form of stochastic optimization problem).

In contrast, there is an alternative approach to solving this problem that we describe as the model-free "black-box" policy optimization approach. In this setting, we forgo learning any model at all of the random variable y. Instead, we simply attempt to learn a policy that maps directly from inputs x to decisions $z^\star(x;\bar{\theta})$ that minimize the loss $L(\theta)$ presented in (2), with the loss defined slightly differently such that $\bar{\theta}$ defines the form of the policy itself (as opposed to defining a predictive model). While such model-free methods can perform well in many settings and can often outperform model-based approaches in settings with vast amounts of data, they often are very data inefficient, as the policy class must have enough representational power to describe sufficiently complex policies without recourse to any underlying model.[3]

Our approach instead offers an intermediate setting, where we *do* still use a model as a surrogate step to determine an optimal decision $z^\star(x;\theta)$, yet we adapt this model specifically based upon the task loss instead of any model prediction accuracy. In practice, we typically want to ultimately minimize some weighted combination of log likelihood *and* the task loss, and the methodology we present here allows us to easily accomplish this.

3.2. Optimizing task loss

To solve the generic optimization problem (3), we can in principle adopt a straightforward (constrained) stochastic gradient approach, as detailed in Algorithm 1. That is, at each iteration we solve the proxy stochastic programming problem (1) to obtain $z^*(x,\theta)$, using the distribution defined by our current values of θ. After doing so, we compute the true loss $L(\theta)$ using the observed value of y. If any of the inequality constraints g_i in $L(\theta)$ are violated, we take a gradient step in this violated constraint; otherwise, we take a step in the gradient of the optimization objective f.

Algorithm 1 Task Loss Optimization

Input: Example $x, y \sim D$.
initialize θ *// some initial parameterization*
for $t = 1, \ldots, T$ **do**
  compute $z^\star(x;\theta)$ via Equation (1)
  initialize α_t *// step size*
  // gradient step in violated constraint or objective
  if $\exists i$ s.t. $g_i(x, y, z^\star(x;\theta)) > 0$ **then**
    update $\theta \leftarrow \theta - \alpha_t \nabla_\theta g_i(x, y, z^\star(x;\theta))$
  else
    update $\theta \leftarrow \theta - \alpha_t \nabla_\theta f(x, y, z^\star(x;\theta))$
  end if
end for

Alternatively, we could move constraints into the objective by adding some appropriate penalty times the positive part of the function, i.e., $\lambda g_i(x, y, z)_+$ for some $\lambda > 0$. In practice, this has the effect of taking gradient steps jointly in all the violated constraints and the objective in the case that one or more inequality constraints are violated.

3.3. Differentiating the optimization solution to a stochastic programming problem

While the above presentation highlights the simplicity of the proposed approach, it avoids the issue of the chief technical challenge to this approach, which is computing the gradient of an objective that depends upon the argmin operation $z^\star(x;\theta)$. Specifically, we need to compute the term

$$\frac{\partial L}{\partial \theta} = \frac{\partial L}{\partial z^\star} \frac{\partial z^\star}{\partial \theta} \tag{5}$$

which involves the Jacobian $\frac{\partial z^\star}{\partial \theta}$. This is the Jacobian of the optimal solution with respect to the distribution parameters θ. Recent approaches have looked into similar argmin differentiations (Gould et al., 2016; Amos et al., 2016), though the methodology we present here is more general and handles the stochasticity of the objective.

We begin by writing the KKT optimality conditions of the general stochastic programming problem (1), where all expectations are taken with respect to the modeled distribution $y \sim p(y|x;\theta)$, which we denote $\mathbf{E}_\theta$; further, assuming the problem is convex means we can replace the general equality constraints $h(x) = 0$ with the linear constraint $Ax = b$. A point (z, λ, ν) is a primal-dual optimal point if

[3] This distinction here is roughly analogous to the policy search vs. model-based settings in reinforcement learning. Indeed, in this context our work bears some similarity to recent work such as value iteration networks, which build models tuned toward reinforcement learning objectives. However, for the purposes of this paper, we typically consider much simpler stochastic programs without the multiple rounds that occur in reinforcement learning, and the extension of these techniques to a full RL setting remains as future work.

it satisfies

$$\begin{aligned} \mathbf{E}_\theta g(z) &\leq 0 \\ Az &= b \\ \lambda &\geq 0 \\ \lambda \circ \mathbf{E}_\theta g(z) &= 0 \\ \nabla_z \mathbf{E}_\theta f(z) + \lambda^T \nabla_z \mathbf{E}_\theta g(z) + A^T \nu &= 0 \end{aligned} \tag{6}$$

where here g denotes the vector of all inequality constraints (represented as a vector-valued function), and where we wrap the dependence on x and y into the functions f and g_i themselves.

Differentiating these equations and applying the implicit function theorem gives a set of linear equations that we can solve to obtain the necessary Jacobians,

$$\begin{bmatrix} \nabla_z^2 \mathbf{E}_\theta f(z) + \sum_{i=1}^{n_{ineq}} \lambda_i \nabla_z^2 \mathbf{E}_\theta g_i(z) & \left(\frac{\partial \mathbf{E}_\theta g}{\partial z}\right)^T & A^T \\ \mathrm{diag}(\lambda)\left(\frac{\partial \mathbf{E}_\theta g}{\partial z}\right) & \mathrm{diag}(\mathbf{E}_\theta g(z)) & 0 \\ A & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{\partial z}{\partial \theta} \\ \frac{\partial \lambda}{\partial \theta} \\ \frac{\partial \nu}{\partial \theta} \end{bmatrix}$$
$$= \begin{bmatrix} \frac{\partial \nabla_z \mathbf{E}_\theta f(z)}{\partial \theta} + \frac{\partial \sum_{i=1}^{n_{ineq}} \lambda_i \nabla_z \mathbf{E}_\theta g_i(z)}{\partial \theta} \\ \mathrm{diag}(\lambda)\frac{\partial g}{\partial z} \\ 0 \end{bmatrix}. \tag{7}$$

The terms in these equations look somewhat complex, but fundamentally, the terms in the left-hand-side linear expressions are the optimality conditions of the convex problem, and the terms on the right hand side are the derivatives of the relevant functions, at the achieved solution, with respect to the governing parameter θ. These equations will take slightly different forms depending on how the stochastic programming problem is solved, but are usually fairly straightforward to compute if the solution is solved in some "exact" manner (i.e., where second order information is used). In practice, we employ sequential quadratic programming (Boggs & Tolle, 1995) to solve these problems, and use a recently-proposed approach for fast solution of the argmin differentiation for QPs (Amos & Kolter, 2017) to solve these linear equations.

4. Experiments

Here we consider two applications of our task-based methodology, one to a synthetic inventory stock problem, and one to an energy scheduling task based on over eight years of real electrical grid data. In both cases, we demonstrate that the task-based end-to-end approach can substantially improve upon other alternatives. Source code for all experiments is available at https://github.com/locuslab/e2e-model-learning.

4.1. Inventory Stock Problem

Problem definition To begin, and to highlight the performance of the algorithm in a setting where the true underlying model is known to us, we consider a "conditional" variation of the classical inventory stock problem (Shapiro & Philpott, 2007). In this setting, a company must decide to order some quantity z of a product in order to minimize costs over some stochastic demand y, whose distribution in turn is affected by some observed features x. There are linear and quadratic costs on product ordered, plus different linear/quadratic costs on overorders $[z-y]_+$ and underorders $[y-z]_+$. That is, the objective is given by:

$$\begin{aligned} f_{stock}(y,z) = c_0 z + \frac{1}{2} q_0 z^2 + c_b[y-z]_+ + \frac{1}{2} q_b([y-z]_+)^2 \\ + c_h[z-y]_+ + \frac{1}{2} q_h([z-y]_+)^2 \end{aligned} \tag{8}$$

where $[v]_+ = \max\{v, 0\}$. Our stochastic programming problem can then be written as

$$\underset{z}{\text{minimize}} \quad \mathbf{E}_\theta[f_{stock}(y,z)]. \tag{9}$$

(Recall that, as y is unknown, we must learn parameters θ that allow us to predict y given inputs x, i.e., the argument of the expectation has a dependence on θ through y.) To simplify the setting, we further assume that the demands are discrete, taking on values $d_1, \ldots, d_k$ with probabilities (conditional on x) $(p_\theta)_i \equiv p(y = d_i | x; \theta)$. Thus our stochastic programming problem (9) can be written succinctly as a joint quadratic program[4]

$$\begin{aligned} \underset{z \in \mathbb{R}, z_b, z_h \in \mathbb{R}^k}{\text{minimize}} \quad & c_0 z + \frac{1}{2} q_0 z^2 + \\ & \sum_{i=1}^k (p_\theta)_i \left(c_b (z_b)_i + \frac{1}{2} q_b (z_b)_i^2 + c_h (z_h)_i + \frac{1}{2} q_h (z_h)_i^2 \right) \\ \text{subject to} \quad & d - z\vec{1} \leq z_b, \quad z\vec{1} - d \leq z_h \\ & z, z_h, z_b \geq 0. \end{aligned} \tag{10}$$

Computation To demonstrate the explicit formula for argmin operation Jacobians for this particular case (e.g., to compute the terms in (7)), note that we can write the above QP in inequality form as $\text{minimize}_{\{\vec{z}: G\vec{z} \leq h\}} \frac{1}{2}\vec{z}^T Q \vec{z} + c^T \vec{z}$

[4] This is often referred to as a two-stage stochastic programming problem (though an extremely trivial example of one), where the first stage variables consist of the amount of product to buy before observing the demand, and the second-stage variables consist of how much to sell back or additionally purchase once the true demand has been revealed.

with

$$\vec{z} = \begin{bmatrix} z \\ z_b \\ z_h \end{bmatrix}, \; Q = \begin{bmatrix} q_0 & 0 & 0 \\ 0 & q_b p_\theta & 0 \\ 0 & 0 & q_h p_\theta \end{bmatrix}, \; c = \begin{bmatrix} c_0 \\ c_b \\ c_h \end{bmatrix}$$
$$G = \begin{bmatrix} -1 & -I & 0 \\ 1 & 0 & -I \\ -1 & 0 & 0 \\ 0 & -I & 0 \\ 0 & 0 & -I \end{bmatrix}, \; h = \begin{bmatrix} -d \\ d \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{11}$$

Thus, for an optimal primal-dual solution $(\vec{z}^\star, \lambda^\star)$, we can compute the Jacobian $\frac{\partial \vec{z}^\star}{\partial p_\theta}$ (the Jacobian of the optimal solution with respect to the probability vector p_θ mentioned above), via the formula

$$\begin{bmatrix} \frac{\partial \vec{z}^\star}{\partial p_\theta} \\ \frac{\partial \lambda^\star}{\partial p_\theta} \end{bmatrix} = \begin{bmatrix} Q & G^T \\ D(\lambda^\star)G & D(G\vec{z}^\star - h) \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ q_b z_b^\star + c_b \vec{1} \\ q_h z_h^\star + c_h \vec{1} \\ 0 \end{bmatrix}, \tag{12}$$

where $D(\cdot)$ denotes a diagonal matrix for an input vector. After solving the problem and computing these Jacobians, we can compute the overall gradient with respect to the task lost $L(\theta)$ via the chain rule

$$\frac{\partial L}{\partial \theta} = \frac{\partial L}{\partial z^\star} \frac{\partial z^\star}{\partial p_\theta} \frac{\partial p_\theta}{\partial \theta} \tag{13}$$

where $\frac{\partial p_\theta}{\partial \theta}$ denotes the Jacobian of the model probabilities with respect to its parameters, which are computed in the typical manner. Note that in practice, these Jacobians need not be computed explicitly, but can be computed efficiently via backpropagation. To both solve the optimization problem in QP form and compute its derivatives, we use a newly-developed differentiable batch QP solver that is described in concurrent work (Amos & Kolter, 2017).

Experimental setup We examine two main cases in our synthetic example to highlight the nature of the algorithm under different conditions: the realizable case, and the non-realizable case.

In all cases we generate problem instances by randomly sampling some $x \in \mathbb{R}^n$, and then generating $p(y|x;\theta)$ according to either $p(y|x;\theta) \propto \exp(\Theta^T x)$ for some $\Theta \in \mathbb{R}^{n \times k}$ (the "realizable" case) or $p(y|x;\theta) \propto \exp((\Theta^T x)^2)$ (the "unrealizable case"). We compare the following approaches on these tasks: 1) An MLE approach that fits a linear model to the data, then computes the allocation by solving the QP; 2) the QP allocation based upon the true model (which gives optimal performance); 3) an end-to-end (nonlinear) neural network policy model; 4) our task-based learning using a linear probability model; and 5) our task-based approach using a non-linear probability model.[5] In all cases we evaluate test performance by running on 1000 random examples, and evaluate performance over 10 folds of different true $\theta^\star$ parameters.

Figure 1(a) shows the performance of these methods in the realizable case. As expected, in this setting the MLE approach performs best, as the true underlying model is in the class of distributions that it can represent and thus solving the stochastic programming problem is a very strong proxy for solving the true optimization problem under the real distribution. Even here, though, the task-based approach substantially outperforms the end-to-end neural network, highlighting the fact that it is more data-efficient to run the learning process "through" a reasonable model. Note that here it does not make a difference whether we use the linear or nonlinear model in the task-based approach.

Figure 1(b) shows the performance on the non-realizable case. Here, the linear MLE, as expected, performs very poorly: it cannot capture the true underlying distribution, and thus solving a stochastic program with this problem would not be expected to perform well. But importantly, the task-based approach with the *linear* model performs much better here: despite the fact that it still has a misspecified model, the task-based nature of the learning process lets us learn a *different* non-linear model than the MLE version, which is particularly tuned to the distribution and loss of the task. Finally, also as to be expected, the non-linear models perform better in this scenario, but again with the task-based non-linear model outperforming the end-to-end policy approach.

4.2. Load Forecasting and Generator Scheduling

We next consider a more realistic grid-scheduling task, based upon over eight years of real electrical grid data. In this setting, a power system operator must decide how much electricity generation $z \in \mathbb{R}^{24}$ to schedule for each hour in the next 24 hours based on some (unknown) distribution over electricity demand. Given a particular realization y of demand, if the operator over-schedules generation ($z > y$), a cost of γ_e per excess unit of production is incurred. In the case of a generation shortage ($z < y$), a cost of γ_s per unit of shortage is incurred. In this setting, $\gamma_s >> \gamma_e$. We also add a quadratic regularization term, in-

[5] We don't explicitly compare to a linear E2E network, as this approach expectedly performs very badly. We also don't compare to a non-linear MLE model, as it would be capable of representing both true underlying models, and the goal here is to demonstrate the algorithm's performance in the case where the true model cannot be realized. In Section 4.2, we highlight a non-linear MLE-based approach on a real-world problem.

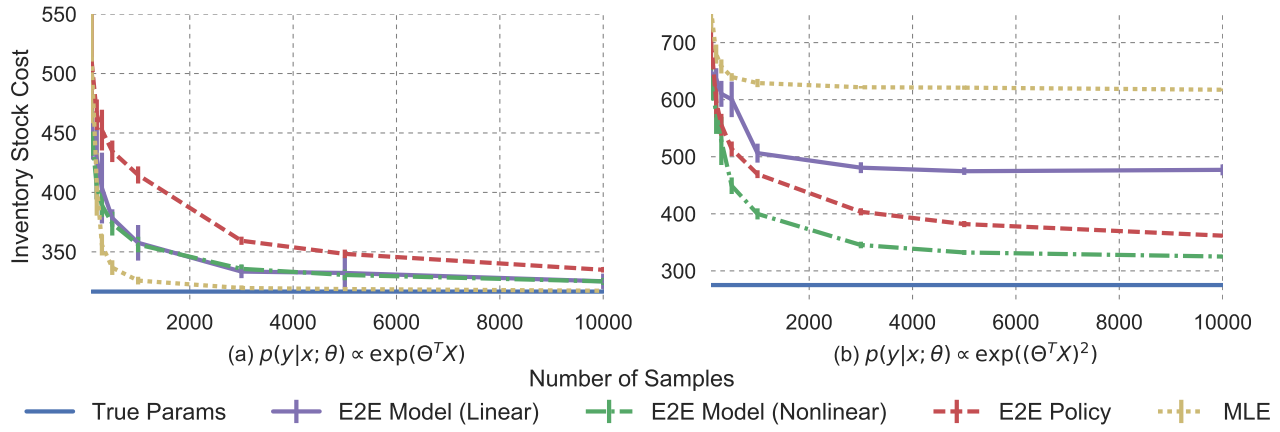

Figure 1. Experimental results for the inventory stock problem in the (a) realizable and (b) non-realizable cases for a representative instantiation of true parameters. (Lower cost is better.) As expected, the MLE approach performs best in the realizable case, but even here, the task-based approach substantially outperforms the end-to-end policy approach. In the non-realizable case, the linear MLE (as expected) performs very poorly, and the task-based approach with the *linear* model performs much better here (despite its misspecified model). The non-linear models both perform better in this scenario, but again with the task-based non-linear model outperforming the end-to-end policy approach.

dicating a preference for generation schedules that closely match demand realizations. Finally, we impose a ramping constraint c_r restricting the difference between the amount of generation scheduled in consecutive timepoints, reflecting the physical limitations associated with quick changes in electricity output levels. These are reasonable proxies for the actual economic costs incurred by electrical grid operators when scheduling generation, and can be written as the stochastic programming problem

$$\underset{z}{\text{minimize}} \quad \sum_{i=1}^{24} \mathbf{E}\left[\gamma_s[y_i - z_i]_+ + \gamma_e[z_i - y_i]_+ + \frac{1}{2}(z_i - y_i)^2\right] \tag{14}$$

where the expectation is taken over the random variable y. Assuming (as we will in our model), that y_i is a Gaussian random variable with mean μ_i and variance σ_i^2, then this expectation has a closed form that can be computed via analytically integrating the Gaussian pdf. Specifically,

$$\begin{aligned}\mathbf{E}\left[\gamma_s[y_i - z_i]_+ + \gamma_e[z_i - y_i]_+ + \frac{1}{2}(z_i - y_i)^2\right] = \\ (\gamma_s + \gamma_e)(\sigma^2 p(z;\mu,\sigma^2) + (z-\mu)F(z;\mu,\sigma^2)) \\ -\gamma_s(z-\mu) + \frac{1}{2}((z_i-\mu_i)^2 + \sigma_i^2)\end{aligned} \tag{15}$$

where $p(z;\mu,\sigma^2)$ and $F(z;\mu,\sigma^2)$ denote the Gaussian pdf and cdf respectively with the given mean and variance. This is a convex function of z (not apparent in this form, but readily established because it is an expectation of a convex function), and we can thus optimize it efficiently and compute the necessary Jacobians. We omit the details due to space constraints, but briefly, we use sequential quadratic programming (SQP) to iteratively approximate the convex objective as a quadratic objective, iterate until convergence, and then compute the necessary Jacobians using the quadratic approximation at the solution, which gives all the correct Hessian and gradient terms. We can furthermore differentiate the gradient and Hessian with respect to the underlying model parameters μ and σ^2 to determine the ultimate gradient of the task loss with respect to the parameters of the distribution.

To develop a predictive model upon which to base our scheduling, we make use of a highly-tuned load forecasting methodology. Specifically, as input to our predictive model, we use the past day's electrical load, the past day's temperature, the future temperature forecast of the next day, as well as additional features such as non-linear functions of the temperatures, binary indicators of weekends or holidays, and yearly sinusoidal features. We then predict the electrical load over all 24 hours of the next day. We use a 2-hidden layer neural network to accomplish this, with an additional residual connection from the inputs to the outputs initialized to the linear regression solution (which can itself perform quite well, though the neural network improves accuracy over the linear model by over 25%). An illustration of the architecture is shown in Figure 2. We train the model to minimize the mean squared error between its predictions and the actual load (this gives the mean prediction μ_i), and compute σ_i^2 as the (constant) empirical variance between the prediction and the actual. In all cases we use 7 years to train the model, and 1.75 subsequent years for testing.

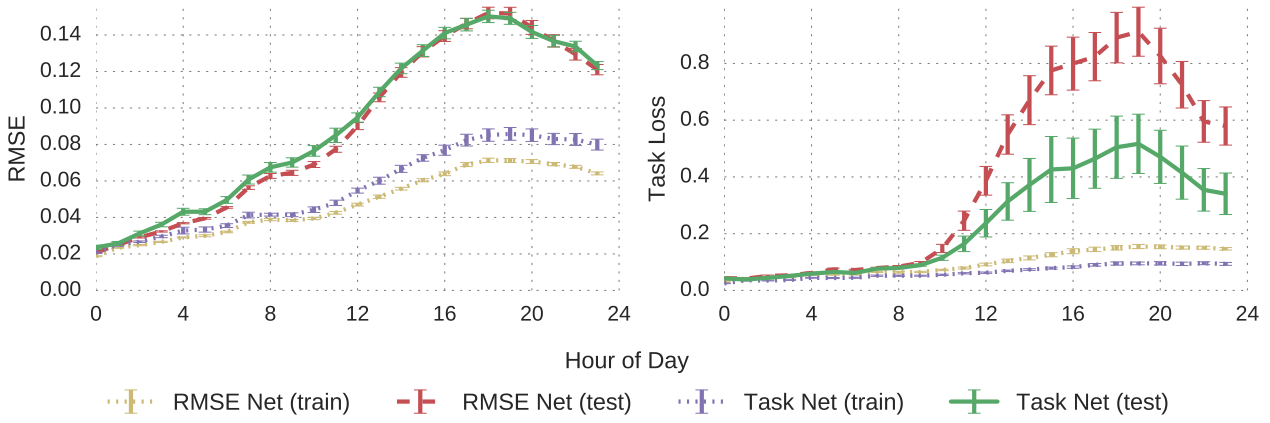

Figure 3. Performance of our "task net" vs. an RMSE-minimizing "RMSE net" on the generation-scheduling problem (lower loss is better). As expected, the RMSE net performs better when measured by the RMSE of its predictions. However, the task net substantially outperforms the RMSE net when evaluated on the task loss (despite the fact that the RMSE network has been highly optimized): in relative terms, we improve upon the performance of the traditional stochastic programming methodology by 25.3%.

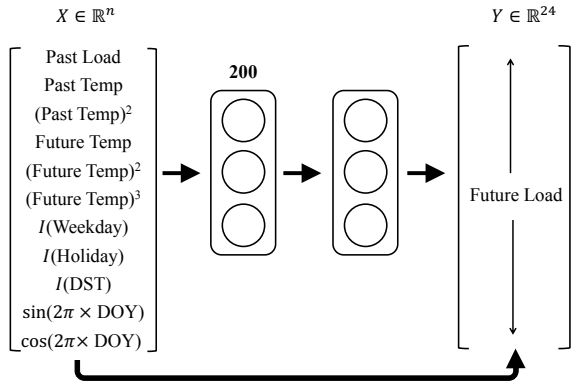

Figure 2. Architecture for the predictive model upon which we base our generation scheduling, mapping from inputs to electrical load predictions over all 24 hours of the next day. We use a 2-hidden layer neural network, with an additional residual connection from the inputs to the outputs initialized to the linear regression solution.

Part of the philosophy behind applying our approach to this model is that we *know* that in reality the Gaussian assumption behind our model is incorrect: the true underlying load follows (precisely) neither a Gaussian distribution nor the common homoskedastic modeling setting where the variance of the distribution is assumed not to depend on the underlying input. However, these assumptions are both exceedingly common in practice, both for the sake of learning models easily *and* for the sake of being able to compute the expectations in the stochastic programming problem exactly. Thus, training the (still Gaussian) systems with a task-based loss has the obvious advantage of retaining computational tractability while simultaneously still allowing us to modify the parameters of the distribution to best improve the actual performance of the task at hand.

Using the predictions of this model (both the mean and variance), we solve the generator scheduling problem by solving the optimization problem (14). In the first case, we learn the parameters of the network by minimizing just the MSE, and in the second case we learn the the parameters by instead minimizing the task loss. A pure end-to-end network is not shown, as it was not able to sufficiently learn the ramp constraints, thus generating infeasible schedules, and we could not obtain good performance even ignoring this infeasibility. Figure 3 shows the performance of the two models, which we refer to as the "RMSE net" and the "task net." As expected, the RMSE net performs better when measured by the RMSE of its predictions; because the load forecasting network is optimized with this objective in mind, it is not surprising that it performs better. However, the task net substantially outperforms the RMSE net when evaluated on the task loss, the actual task objective that the system operator cares about (despite the fact that the RMSE network has been highly optimized): in relative terms, we improve upon the performance of the traditional stochastic programming methodology by 25.3%.

5. Conclusions and Future Work

In this paper, we propose an end-to-end approach for learning probabilistic machine learning models that will eventually be used in the loop of a larger process. Specifically, we consider training probabilistic models in the context of stochastic programming, in a manner that directly captures the ultimate task-based objective for which they will be used. Given the reality that all models make errors,

the proposed method allows the modeler to make explicit tradeoffs between model error over different regions of the input/output space, which may be more appropriate in the specified context than the implicit tradeoffs made in a pure maximum log likelihood setting. On the flip side, our end-to-end model learning approach is potentially less data-intensive than a pure end-to-end policy learning approach, as the latter approach requires many historical training examples to compensate for the lack of recourse to any underlying statistical model.

Preliminary experiments indicate that our proposed task-based learning approach outperforms pure end-to-end policy learning approaches in the realizable case where the true probability distribution is contained within our model class), and outperforms an MLE approach in the non-realizable case. Our approach also achieves a 25.3% performance improvement over a traditional real-world stochastic programming methodology for scheduling electricity generation based on forecasted electricity load, despite the highly-optimized nature of the traditional model. As such, our approach provides a promising methodology for optimizing probabilistic machine learning models that will then be used to optimize some other process.

Future work includes extension of our approach to stochastic learning models with multiple rounds, and further to a full reinforcement learning setting. Future work may also include incorporating an informed search over various model classes in the context of our proposed hybrid approach, in order to improve the chances of task-consistency and thereby of better model performance.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1252522.

References

Amodei, Dario, Anubhai, Rishita, Battenberg, Eric, Case, Carl, Casper, Jared, Catanzaro, Bryan, Chen, Jingdong, Chrzanowski, Mike, Coates, Adam, Diamos, Greg, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. *arXiv preprint arXiv:1512.02595*, 2015.

Amos, Brandon and Kolter, J Zico. Optnet: Differentiable optimization as a layer in neural networks. *arXiV preprint arXiv:1703.00443*, 2017.

Amos, Brandon, Xu, Lei, and Kolter, J Zico. Input convex neural networks. *arXiv preprint arXiv:1609.07152*, 2016.

Bengio, Yoshua. Using a financial training criterion rather than a prediction criterion. *International Journal of Neural Systems*, 8(04):433–443, 1997.

Boggs, Paul T and Tolle, Jon W. Sequential quadratic programming. *Acta numerica*, 4:1–51, 1995.

Buzacott, John A and Shanthikumar, J George. *Stochastic models of manufacturing systems*, volume 4. Prentice Hall Englewood Cliffs, NJ, 1993.

Gould, Stephen, Fernando, Basura, Cherian, Anoop, Anderson, Peter, Cruz, Rodrigo Santa, and Guo, Edison. On differentiating parameterized argmin and argmax problems with application to bi-level optimization. *arXiv preprint arXiv:1607.05447*, 2016.

Graves, Alex and Jaitly, Navdeep. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*, volume 14, pp. 1764–1772, 2014.

Hazan, Tamir, Keshet, Joseph, and McAllester, David A. Direct loss minimization for structured prediction. In *Advances in Neural Information Processing Systems*, pp. 1594–1602, 2010.

He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, and Sun, Jian. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

LeCun, Yann, Muller, Urs, Ben, Jan, Cosatto, Eric, and Flepp, Beat. Off-road obstacle avoidance through end-to-end learning. In *NIPS*, pp. 739–746, 2005.

Levine, Sergey, Finn, Chelsea, Darrell, Trevor, and Abbeel, Pieter. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39): 1–40, 2016.

Linderoth, Jeff, Shapiro, Alexander, and Wright, Stephen. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, 142 (1):215–241, 2006.

Ng, Andrew Y and Jordan, Michael I. On discriminative vs. generative classifiers: A comparison of logistic regression and naive bayes. *Advances in neural information processing systems*, 2:841–848, 2002.

Rockafellar, R Tyrrell and Wets, Roger J-B. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of operations research*, 16(1):119–147, 1991.

Shapiro, Alexander and Philpott, Andy. A tutorial on stochastic programming. *Manuscript. Available at www2. isye. gatech. edu/ashapiro/publications. html*, 17, 2007.

Song, Yang, Schwing, Alexander G, Zemel, Richard S, and Urtasun, Raquel. Training deep neural networks via direct loss minimization. In *Proceedings of The 33rd International Conference on Machine Learning*, pp. 2169–2177, 2016.

Stoyanov, Veselin, Ropson, Alexander, and Eisner, Jason. Empirical risk minimization of graphical model parameters given approximate inference, decoding, and model structure. *International Conference on Artificial Intelligence and Statistics*, 15:725–733, 2011. ISSN 15324435. URL http://machinelearning.wustl.edu/mlpapers/paper{_}files/AISTATS2011{_}StoyanovRE11.pdf{%}5Cnpapers2://publication/uuid/605A8634-EE37-4042-A4B5-E41E93EDA7EF.

Tamar, Aviv, Levine, Sergey, Abbeel, Pieter, WU, YI, and Thomas, Garrett. Value iteration networks. In *Advances in Neural Information Processing Systems*, pp. 2146–2154, 2016.

Thomas, Ryan W, Friend, Daniel H, Dasilva, Luiz A, and Mackenzie, Allen B. Cognitive networks: adaptation and learning to achieve end-to-end performance objectives. *IEEE Communications Magazine*, 44(12):51–57, 2006.

Wallace, Stein W and Fleten, Stein-Erik. Stochastic programming models in energy. *Handbooks in operations research and management science*, 10:637–677, 2003.

Wang, Kai, Babenko, Boris, and Belongie, Serge. End-to-end scene text recognition. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pp. 1457–1464. IEEE, 2011.

Wang, Tao, Wu, David J, Coates, Adam, and Ng, Andrew Y. End-to-end text recognition with convolutional neural networks. In *Pattern Recognition (ICPR), 2012 21st International Conference on*, pp. 3304–3308. IEEE, 2012.

Ziemba, William T and Vickson, Raymond G. *Stochastic optimization models in finance*, volume 1. World Scientific, 2006.